

The cover time, the blanket time, and the Matthews bound

J. Kahn *

J. H. Kim †

L. Lovász ‡

V. H. Vu §

February 1, 2008

Abstract

We prove upper and lower bounds and give an approximation algorithm for the cover time of the random walk on a graph. We introduce a parameter M motivated by the well known Matthews bounds on the cover time and prove that $M/2 \leq C = O(M(\ln \ln n)^2)$. We give a deterministic polynomial time algorithm to approximate M within a factor of 2; this then approximates C within a factor of $O((\ln \ln n)^2)$, improving previous bound of $O(\ln n)$ of Matthews.

The blanket time B was introduced by Winkler and Zuckerman: it is the expectation of the first time when all vertices are visited within a constant factor of number of times suggested by the stationary distribution. Obviously $C \leq B$, and they conjectured $B = O(C)$ and proved $B = O(C \ln n)$. Our bounds above are also valid for the blanket time, and so it follows that $B = O(C(\ln \ln n)^2)$.

1 Introduction

Given a connected graph G on n vertices, for a vertex $i \in V(G)$, $C(i)$ denotes the cover time of the usual random walk on G , starting from i ; that is, $C(i)$ is the expectation of the number of steps a random walk starting from i takes until it covers all vertices of G . The quantity $C = \max_{i \in V(G)} C(i)$ is called the cover time of

G . (See [1] for background.)

Although C is a basic notion in the theory of random walks, there is no effective way known to compute this parameter, given the adjacency matrix of G as the input. The following question has been open for several years [1].

Question. *Is there a deterministic algorithm which approximates C up to a constant factor in polynomial time ?*

The requirement that the algorithm is deterministic is crucial and this makes the problem difficult. It is simple to provide a randomized algorithm which approximates C within a factor $(1+\epsilon)$ for any positive constant ϵ , with high probability: just simulate the chain and take the average of the empirical cover times.

Prior to this paper, the best approximation factor we knew of was $\ln n$. This factor can be achieved using the following fundamental result of Matthews [3]. For any pair of vertices $i, j \in V(G)$, $H(i, j)$ denotes the hitting time from i to j . We set

$$h_{\max} = \max_{i, j \in V} H(i, j), \quad h_{\min} = \min_{i, j \in V} H(i, j),$$

and more generally, for every set $S \subseteq V$, we let

$$h_S = \min_{\substack{i, j \in S \\ i \neq j}} H(i, j).$$

$$\text{Let } \text{har}(n) = \sum_{i=1}^n 1/i.$$

Theorem 1.1 (Matthews' theorem) *For any G ,*

$$h_{\min} \text{har}(n) \leq C \leq h_{\max} \text{har}(n).$$

More generally, for any subset $S \subset V(G)$ with $|S| \geq 2$,

$$h_S \text{har}(|S|) \leq C. \quad (1)$$

*Department of Mathematics, Rutgers University, NJ 08903; jkahn@math.rutgers.edu. Supported in part by NSF.

†Microsoft Research, Microsoft Corporation, Redmond, WA 98052; jehkim@microsoft.com

‡Microsoft Research, Microsoft Corporation, Redmond, WA 98052; lovasz@microsoft.com

§Microsoft Research, Microsoft Corporation, Redmond, WA 98052; vanhavu@microsoft.com

It follows from the upper bound in Matthews' theorem and the definition of the cover time that

$$h_{\max} \leq C \leq h_{\max} \text{har}(n).$$

Thus, h_{\max} approximates C within a factor of $\text{har}(n) \approx \ln n$. Moreover, since the $H(i, j)$'s are quite easily computable in polynomial time, h_{\max} is computable in polynomial time. Unfortunately, h_{\max} can be equal to the cover time (as shown by a path), as well as a factor of $\ln n$ off the cover time (as shown by a complete graph).

We could try to use the lower bound (1) in Matthews' Theorem, by maximizing over S . Unfortunately, this can be even worse than h_{\max} . For example, if G consists of a single edge with N loops added at one of the nodes, then the best Matthews lower bound is 1, while the cover time, starting from either the node with the loops or from the stationary distribution, is about N . This problem is easy to fix: just throw in the obvious lower bound h_{\max} . More precisely, let M_0 be the maximum of h_{\max} and the quantities $h_S \ln |S|$ ($S \subseteq V$, $|S| \geq 2$). Then M_0 is a lower bound on C , and (as we'll see) it is only a $(\ln \ln n)^2$ factor off. We call M_0 the *augmented Matthews bound*.

A parameter closely related to the cover time is the *blanket time*, introduced by Winkler and Zuckerman [5]. The definition provided below is a little bit stronger.

Definition. Consider a random walk starting from a node v . Let $r_{T,v}(x)$ be the number of visits to x up to time T . Let \mathbf{B} be the first time T when the ratio $\frac{r_{T(i)}/\pi_i}{r_{T(j)}/\pi_j}$ is at most 2 for any two nodes i and j (in particular, all nodes are covered by this time). Let $B(v)$ be the expectation \mathbf{B} . The blanket time B is the maximum of $B(v)$, over all vertices v .

It is clear that $C \leq B$. Winkler and Zuckerman conjectured that there is a constant K so that $B \leq KC$, and showed that

$$B = O(C \ln n).$$

The main goal of this paper is to improve the factor $O(\ln n)$ in both problems mentioned above to $O((\ln \ln n)^2)$.

The following variant of the augmented Matthews bound M_0 is at the heart of our study.

Let $\kappa(i, j) = H(i, j) + H(j, i)$ be the commute time between i and j . For any $S \subset V(G)$, let $\kappa_S = \min_{i, j \in S} \kappa(i, j)$, and

$$M = \max_{S \subset V(G)} \kappa_S \ln |S|.$$

As the following proposition shows, M and M_0 are essentially equivalent, but due to the symmetry of κ , M will be easier to handle.

Proposition 1.2

$$\frac{1}{8}M \leq M_0 \leq M.$$

Our main theorem is the following.

Theorem 1.3 *For every graph G on n vertices*

$$\frac{1}{2}M \leq C \leq B \leq 10^5 M (\ln \ln n)^2$$

(Of course the lower bound $C \geq M/8$ follows from Proposition 1.2 and Matthews' bound.)

It follows from this theorem that M approximates C within a factor $O((\ln \ln n)^2)$ and $B \leq KC(\ln \ln n)^2$, for some constant K . It turns out, somewhat surprisingly, that both the upper bound and the lower bound are sharp up to a constant factor.

The proof of the lower bound will give a somewhat stronger result. Let $C(\pi) = \sum_i \pi_i C(i)$ denote the cover time when the walk is started from a random node from the stationary distribution π .

Theorem 1.4 *For any graph G ,*

$$C(\pi) \geq \frac{1}{2}M.$$

An important property of M as an approximation of the cover time is that it is efficiently approximable:

Theorem 1.5 *M can be approximated within a factor of 2 by a deterministic polynomial algorithm.*

The rest of the paper is divided into five sections. In Section 2, we describe an algorithm which computes M up to a factor of 2, proving Theorem 1.5. In Section 3, as preparation for the proof of the lower bound in Theorem 1.3, we derive some formulas for the cover time, which may be interesting in their own right. In Section 4, we complete the proof of Theorem 1.4, and also prove Proposition 1.2. The proof of the upper bound in Theorem 1.3, which is the most substantial part of this paper, follows in Section 5. In the final Section 6, we give constructions which show that both the upper and lower bounds in Theorem 1.3 can be attained.

2 Approximating M

Since the commute times $\kappa(i, j)$ are polynomially computable, the quantity κ_S is also polynomially computable for any set $S \subset V(G)$. However, the definition of M involves all (exponentially many) subsets of $V(G)$ and it is not clear that one can compute M in polynomial time. In the following, we show that one can, at least, approximate M to within a factor of 2 in polynomial time.

A preliminary remark: the commute time $\kappa(i, j)$ satisfies the triangle inequality:

$$\kappa(i, k) \leq \kappa(i, j) + \kappa(j, k),$$

and hence we can consider it as a “distance” on the graph.

Algorithm. To start, pick an arbitrary vertex v_1 . At the i^{th} step ($i = 1, 2, \dots, n$), we have selected the set $V_i = \{v_1, \dots, v_i\}$. Choose v_{i+1} to be a vertex $v \in V \setminus V_i$ whose distance $\min_{u \in V_i} \kappa(u, v)$ from V_i is maximum. Compute $M_i = \kappa_{V_i} \ln i$ for all $i = 2, 3, \dots, n$, and output $M' = \max_i M_i$.

Since $\kappa(i, j)$ are polynomially computable, our algorithm runs in polynomial time. Moreover, $M' \leq M$ by definition. It remains to show that $2M' \geq M$.

Assume that M is attained at a set $S \subset V(G)$ of cardinality s . We claim that $2M_s \geq M$. It suffices to show that $\kappa_{V_s} \geq \kappa_S/2$.

Let $R = S \setminus V_s$. If R is empty then $S = V_s$ and we are done, so we assume that $|R| = r > 0$. By

the description of the algorithm, $\kappa_{V_s} = \kappa(v_s, v_j)$ for some $j < s$.

For each vertex $x \in R$, there is a vertex $y_x \in V_{s-1}$ so that $\kappa(x, y_x) \leq \kappa_{V_s}$. If $y_x \in S$ for some x , then $\kappa_S \leq \kappa(x, y_x) \leq \kappa_{V_s}$, and we are done. If $y_x \in V_{s-1} \setminus S$ for all $x \in R$, then (using that $|V_{s-1} \setminus S| = (s-1) - (s-r) = r-1 < r$) the pigeon hole principle gives that there are x and x' in R so that $y_x = y_{x'} = y$. So by the triangle inequality

$$\kappa(x, x') \leq \kappa(x, y) + \kappa(y, x') \leq 2\kappa_{V_s}.$$

By definition $\kappa(x, x') \geq \kappa_S$ and the proof is complete.

Remark. The only property of the commute times we use here is the triangle inequality. Therefore, our result holds in a more general setting. Consider a metric w on a finite set V of n points. For any subset $S \subset V$, let $w_S = \min_{i, j \in S} w(i, j)$ (if S has less than 2 elements, $w_S = 0$). Define

$$W = \max_{S \subset V} w_S f(|S|),$$

where f is any non-negative function defined on the set of non-negative integers.

Corollary 2.1 *For any finite metric space and any non-negative f , the above algorithm (with $\kappa(i, j)$ replaced by w_{ij}) computes W within a factor of 2.*

3 Formulas for the cover time

Fix a set $S \subseteq V$, $|S| = s \geq 2$, and a starting node v . For a given random walk ($v = v^0, v^1, v^2, \dots$), and a set $T \subseteq S$, let $Z(T)$ denote the set of nodes of S not seen before T is first reached. Thus $T \subseteq Z(T)$. Define, for $i, j \in S$,

$$A(i, j) = \begin{cases} \frac{1}{|Z(i)|(|Z(i)|-1)}, & \text{if } j \in Z(i) \setminus \{i\}, \\ 0, & \text{otherwise,} \end{cases}$$

(this number depends on the walk) and let $a(i, j) = \mathbb{E}[A(i, j)]$. We have

$$\begin{aligned} \sum_{i \in S} \sum_{j \in S} A(i, j) &= \sum_{i \in S} \sum_{j \in Z(i) \setminus \{i\}} \frac{1}{|Z(i)|(|Z(i)|-1)} \\ &= \sum_{i: |Z(i)| > 1} \frac{1}{|Z(i)|} = \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{s}, \end{aligned}$$

and thus

$$\sum_{i \in S} \sum_{j \in S} a(i, j) = \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{s} \approx \ln s. \quad (2)$$

Using this notation, we can state a formula for the expected number $C(v, S)$ of steps until all nodes of S are visited. The basic idea here is similar to that in Mathhews' theorem.

Lemma 3.1

$$C(v, S) = \frac{1}{s} \sum_{j \in S} H(v, j) + \sum_{i \in S} \sum_{j \in S} H(i, j) a(i, j).$$

Proof. Let X_k be the number of steps required to see k nodes of S . Clearly $C(v, S) = \mathbb{E}[X_s]$. Let $T(i)$ be the number of steps required to see node i . The following algebraic identity is easy to verify:

$$\begin{aligned} X_s &= \frac{1}{s} \sum_{k=1}^s X_k \\ &+ \sum_{1 \leq k < m \leq s} \frac{1}{(s-k)(s-k+1)} (X_m - X_k) \end{aligned} \quad (3)$$

Now here

$$\mathbb{E} \left[\sum_{k=1}^s X_k \right] = \sum_{i \in S} \mathbb{E}[T(i)] = \sum_{i \in S} H(v, i).$$

For the second sum in (3), we fix the first X_k steps, then

$$\sum_{m: k \leq m \leq s} (X_m - X_k) = \sum_{j \in Z(v^{X_k})} (T(j) - T(v^{X_k})),$$

and hence

$$\mathbb{E} \left[\sum_{m: k \leq m \leq s} (X_m - X_k) \right] = \sum_{j \in Z(v^{X_k})} H(v^{X_k}, j).$$

Summing over k , we get

$$\begin{aligned} &\sum_{k=1}^{s-1} \frac{1}{(s-k)(s-k+1)} \sum_{j \in Z(v^{X_k})} H(v^{X_k}, j) \\ &= \sum_{i \in S} \frac{1}{(|Z(i)|-1)|Z(i)|} \sum_{j \in Z(i)} H(i, j) \end{aligned}$$

$$= \sum_{i \in S} \sum_{j \in S} H(i, j) A(i, j).$$

Taking expectation again, we get the lemma. \square

For $i, j \in S$, let

$$Q(i, j) = \frac{1}{|Z(ij)|(|Z(ij)|-1)},$$

and $q(i, j) = a(i, j) + a(j, i) = \mathbb{E}[Q(i, j)]$. Using the identity

$$H(\pi, j) - H(\pi, i) = H(i, j) - H(j, i). \quad (4)$$

due to Tetali and Winkler [4], which implies that

$$H(i, j) = \frac{1}{2} \kappa(i, j) + \frac{1}{2} (H(\pi, j) - H(\pi, i)), \quad (5)$$

a simple computation gives the following lemma:

Lemma 3.2

$$C(v, S) = \frac{1}{s} \sum_{j \in S} (H(v, j) - H(\pi, j)) \quad (6)$$

$$+ \frac{1}{4} \sum_{i \in S} \sum_{j \in S} \kappa(i, j) q(i, j). \quad (7)$$

Let $q_\pi(i, j)$ be the expectation of $q(i, j)$, when the starting node v is chosen at random from the stationary distribution. Averaging over v , the first term in (6) cancels, and we get

Corollary 3.3

$$C(\pi, S) = \frac{1}{4} \sum_{i \in S} \sum_{j \in S} \kappa(i, j) q_\pi(i, j).$$

4 Proof of the lower bound.

Proof of Theorem 1.4 This follows easily from Corollary 3.3 and (2):

$$C(\pi) \geq C(\pi, S) = \frac{1}{4} \sum_{i \in S} \sum_{j \in S} \kappa(i, j) q_\pi(i, j)$$

$$\geq \frac{1}{4} \kappa_S \sum_{i \in S} \sum_{j \in S} q_\pi(i, j) = \frac{1}{2} \kappa_S \ln |S|.$$

Proof of Proposition 1.2. It is obvious that for every $S \subseteq V$

$$h_S \ln |S| \leq \frac{1}{2} \min_{i, j \in S, i \neq j} \kappa(i, j) \ln |S| \leq \frac{1}{2} M,$$

and for every $i, j \in V$,

$$H(i, j) \leq \kappa(i, j) \leq M.$$

Hence $M_0 \leq M$.

To prove the other bound, let S be the set attaining the maximum in the definition of M . If $h_{\max} > M/4$, then we have nothing to prove, so suppose that $H(i, j) \leq M/4$ for all i and j .

We define a digraph D on S as follows. There is an edge from i to j if and only if $H(i, j) \leq \kappa(i, j)/4$. In this case, it is clear that $H(j, i) - H(i, j) \geq \kappa(i, j)/2 \geq \kappa_S/2$.

Let $i_0 i_1 \dots i_m$ be a (directed) path of length m in D . Then by the cycle law [4] we have

$$H(i_m, i_0) \geq \sum_{l=0}^{m-1} (H(i_{l+1}, i_l) - H(i_l, i_{l+1})) \geq m\kappa_S/2.$$

On the other hand, $H(i, j) \leq \kappa_S \ln |S|/4$ for all $i, j \in S$. This implies that $m \leq \ln |S|/2$. A theorem of Gallai [2] implies that D is $\ln |S|/2$ colorable, and therefore D contains an independent set I of size at least $2|S|/\ln |S| > |S|^{1/2}$. By the definition of D ,

$$H(i, j) \geq \kappa(i, j)/4 \geq \kappa_S/4$$

for any $i, j \in I$. Therefore,

$$\min_{i, j \in I} H(i, j) \ln |I| \geq \frac{1}{4} \kappa_S \frac{1}{2} \ln |S| = \frac{1}{8} M.$$

5 Proof of the upper bound

We need a Chernoff type large deviation inequality, which will be shown using fairly standard arguments.

Lemma 5.1 *Let X_1, \dots, X_k be independent non-negative integer valued random variables with*

$$\mathbb{P}[X_i = m] \leq a(1-p)^m \quad \forall m \geq 1$$

for some numbers $a > 0$ and $0 < p < 1$. Let $X = \sum_{i=1}^k X_i$. Then for any $L > 0$

$$\mathbb{P}[X - \mathbb{E}[X] \leq -L] \leq \exp\left(-\frac{p^3 L^2}{4(1-p)ak}\right).$$

Proof. As usual, we first estimate $\mathbb{E}[e^{-\lambda X_i}]$ for $\lambda > 0$. Taylor expansion gives

$$\begin{aligned} \mathbb{E}[e^{\lambda X_i}] &= 1 - \lambda \mathbb{E}[X_i] + (\lambda^2/2) \mathbb{E}[X_i^2 e^{-\lambda^* X_i}] \\ &\leq \exp(-\lambda \mathbb{E}[X_i] + (\lambda^2/2) \mathbb{E}[X_i^2]) \end{aligned}$$

for some λ^* between 0 and λ . Since

$$\begin{aligned} \mathbb{E}[X_i^2] &= \sum_{m=1}^{\infty} m^2 \mathbb{P}[X_i = m] \\ &\leq a \sum_{m=1}^{\infty} m^2 (1-p)^m \end{aligned}$$

and

$$\sum_{m=1}^{\infty} m^2 (1-p)^m = \frac{(1-p) + (1-p)^2}{p^3} \leq \frac{2(1-p)}{p^3},$$

it follows that

$$\mathbb{E}[e^{-\lambda X_i}] \leq \exp\left(-\lambda \mathbb{E}[X_i] + \frac{a(1-p)\lambda^2}{p^3}\right).$$

Therefore

$$\begin{aligned} \mathbb{E}[e^{-\lambda(X - \mathbb{E}[X])}] &= \prod_i \mathbb{E}[e^{-\lambda(X_i - \mathbb{E}[X_i])}] \\ &\leq \exp\left(\frac{ak(1-p)\lambda^2}{p^3}\right). \end{aligned}$$

□ Taking $\lambda = p^3 L / (2(1-p)ak)$, we have

$$\begin{aligned} \mathbb{P}[X - \mathbb{E}[X] \leq -L] &\leq \mathbb{E}[e^{-\lambda(X - \mathbb{E}[X] + L)}] \\ &\leq \exp\left(-\lambda L + \frac{ak(1-p)\lambda^2}{p^3}\right) \\ &= \exp\left(-\frac{p^3 L^2}{4(1-p)ak}\right). \end{aligned}$$

□

Lemma 5.2 *Let i and j be two nodes and $k \geq 1$. Let W_k be the number of times j had been visited when i was visited the k -th time. Then for every $\varepsilon > 0$,*

$$\mathbb{P}\left[W_k < (1-\varepsilon) \frac{\pi_j}{\pi_i} k\right] \leq \exp\left(\frac{-\varepsilon^2 k}{4\pi_i \kappa(i, j)}\right)$$

Proof. Let us restrict the Markov chain to i and j only. It is well known that we get a time-reversible Markov chain with transition probabilities

$$\begin{aligned}\hat{p}_{ii} &= 1 - \frac{1}{\pi_i \kappa(i, j)} & \hat{p}_{ij} &= \frac{1}{\pi_i \kappa(i, j)} \\ \hat{p}_{ji} &= \frac{1}{\pi_j \kappa(i, j)} & \hat{p}_{jj} &= 1 - \frac{1}{\pi_j \kappa(i, j)}\end{aligned}$$

and stationary probabilities

$$\hat{\pi}_i = \frac{\pi_i}{\pi_i + \pi_j}, \quad \hat{\pi}_j = \frac{\pi_j}{\pi_i + \pi_j}.$$

We may consider this very simple Markov chain to prove the lemma.

Define X_k to be the number of visits to j during the k^{th} return trip from i to itself, that is,

$$X_j = W_{k+1} - W_k.$$

It is clear that the X_j are i.i.d. with

$$\begin{aligned}\mathbb{P}[X_1 = 0] &= \hat{p}_{ii} \\ \mathbb{P}[X_1 = m] &= \hat{p}_{ij} \hat{p}_{jj}^{m-1} \hat{p}_{ji} \\ \mathbb{E}[X_1] &= \frac{\pi_j}{\pi_i}\end{aligned}$$

Clearly $W_k = X_1 + \dots + X_k$. Thus

$$\mathbb{E}[W_k] = k \mathbb{E}[X_1] = k \frac{\pi_j}{\pi_i}.$$

Applying Lemma 5.1 with $a = \hat{p}_{ij} \hat{p}_{ji} (1 - \hat{p}_{ji})^{-1}$, $p = \hat{p}_{ji}$ and $L = \varepsilon \frac{\pi_j}{\pi_i} k$, we obtain

$$\begin{aligned}\mathbb{P}[W_k < (1 - \varepsilon) \frac{\pi_j}{\pi_i} k] &= \mathbb{P}\left[X - \mathbb{E}[X] \leq -\varepsilon \frac{\pi_j}{\pi_i} k\right] \\ &\leq \exp\left(-\frac{\hat{p}_{ji}^3 \varepsilon^2 \pi_i^2 k^2}{4 \pi_j^2 \hat{p}_{ij} \hat{p}_{ji} k}\right) = \exp\left(-\frac{\varepsilon^2 k}{4 \pi_i \kappa(i, j)}\right)\end{aligned}$$

□

Proof of the upper bound in Theorem 1.3. Consider the ordering $(v_0, v_1, \dots, v_{n-1})$ of the nodes of G as obtained by the Algorithm in section 2. For convenience, relabel the nodes by $(1, \dots, n)$. Recall that each $i > 1$ is a node farthest away from the set $\{1, \dots, i-1\}$ in distance κ .

For each node $i > 1$, let i' be a node with $i' \leq \sqrt{i}$ and $\kappa(i, i')$ minimal. Clearly, the edges ii' form a tree \mathcal{T} . We consider 1 as the root of the tree. It is also clear that the depth d of \mathcal{T} is at most $1.5 \ln \ln n$.

Our next observation is that

$$\kappa(i, i') \leq \frac{2M}{\ln i}. \quad (8)$$

Indeed, let $S = \{1, \dots, \lfloor \sqrt{i} \rfloor + 1\}$. Then, by the definition of M ,

$$\kappa_S \leq \frac{M}{\ln |S|} < \frac{2M}{\ln i},$$

and hence there exist nodes $u, v \in S$ with $\kappa(u, v) < 2M/\ln i$. We may assume that $u \leq v$. By the choice of the ordering, there exists a node $j \leq v-1 \leq \sqrt{i}$ such that $\kappa(i, j) \leq \kappa(u, v)$. It follows that

$$\kappa(i, i') \leq \kappa(u, v) \leq \frac{2M}{\ln i}.$$

Set $\varepsilon = 1/(8 \ln \ln n)$ and $T_0 = \lceil 400M/\varepsilon^2 \rceil$. Our next goal is to bound the probability that $\mathbf{B} > T$ for some $T \geq T_0$.

Set $F(i) = r_T(i)/(T \pi_i)$. On the average, $F(i) = 1$. If the event “ $\mathbf{B} > T$ ” occurs, then there exists an edge ii' of \mathcal{T} with one of the following properties:

- (A) $F(i') \geq 0.9(1 + \ln i')^{-\varepsilon}$ and $F(i) < 0.9(1 + \ln i)^{-\varepsilon}$;
- (B) $F(i') \leq 1.1(1 + \ln i')^{\varepsilon}$ and $F(i) > 1.1(1 + \ln i)^{\varepsilon}$;
- (C) $F(i') \leq 0.9(1 + \ln i')^{\varepsilon}$ and $F(i) > 0.9(1 + \ln i)^{\varepsilon}$;
- (D) $F(i') \geq 1.1(1 + \ln i')^{\varepsilon}$ and $F(i) < 1.1(1 + \ln i)^{\varepsilon}$.

Indeed, if \mathbf{B} is larger than T , then there exists a node u such that either $F(u) > \sqrt{2}$ or $F(u) < 1/\sqrt{2}$. Suppose that e.g. the second occurs. We assume that $n > 10$, to exclude some trivial complications. Then $F(u) < 0.9(1 + \ln u)^{-\varepsilon}$. We also know that there is a node v with $F(v) > 1 > 0.9(1 + \ln v)^{\varepsilon}$. If $F(1) > 0.9$, then along the path from u to 1 there is an edge with property (A). If $F(1) \leq 0.9$, then along the path from v to 1 there is an edge with property (C).

We call such an edge “bad”. To bound the probability that an edge is bad, we have to bound the probabilities of (A), (B), (C) and (D) separately. This is very similar in all cases, and we give the details for (A). Let $k = \lceil 0.9(1 + \ln i')^{-\varepsilon} \pi_{i'} T \rceil$, and

consider the step when i' is reached the k -th time. By (A), the number of times we have seen i is

$$\begin{aligned} W_k &< 0.9(1 + \ln i)^{-\varepsilon} \pi_i T \leq \left(\frac{1 + \ln i}{1 + \ln i'} \right)^{-\varepsilon} \frac{\pi_i}{\pi_{i'}} \cdot k \\ &< 2^{-\varepsilon} \frac{\pi_i}{\pi_{i'}} \cdot k < \left(1 - \frac{\varepsilon}{4} \right) \frac{\pi_i}{\pi_{i'}} \cdot k, \end{aligned}$$

and hence by Lemma 5.2,

$$\mathbb{P}[A] \leq \exp \left(\frac{-\varepsilon^2 k}{100 \pi_{i'} \kappa(i, i')} \right).$$

Now here

$$\frac{k}{\pi_{i'}} \geq 0.9(1 + \ln i')^{-\varepsilon} T \geq \frac{1}{2} T$$

and hence by (8),

$$\mathbb{P}[A] < \exp \left(\frac{-\varepsilon^2 T \ln i}{200M} \right) < i^{-T\varepsilon^2/(200M)}.$$

The probability that this happens for some edge ii' is at most

$$\sum_{i=2}^n i^{-T\varepsilon^2/(200M)} < 2 \cdot 2^{-T\varepsilon^2/(200M)},$$

(using here that $T \geq T_0$) and hence the probability that $\mathbf{B} > T$ is at most $8 \cdot 2^{-T\varepsilon^2/(200M)}$. Thus

$$\begin{aligned} \mathbb{E}[\mathbf{B}] &= \sum_{T=0}^{\infty} \mathbb{P}[\mathbf{B} > T] \leq T_0 + \sum_{T=T_0}^{\infty} \mathbb{P}[\mathbf{B} > T] \\ &\leq T_0 + 8 \sum_{T=T_0}^{\infty} 2^{-T\varepsilon^2/(200M)} < 2T_0, \end{aligned}$$

which proves the theorem. \square

Remark. We may prove the upper bound using a slightly different approach. Let S_0 be the set of all vertices and inductively define S_i to be a maximal subset of S_{i-1} such that $\kappa_{S_i} > \kappa_i := 2^i M / \log n$. If such a subset does not exist, S_i consists of a vertex in S_{i-1} and the construction stops. Since $\kappa_{S_i} \leq M$ unless $|S_i| = 1$, this procedure stops within $O(\ln \ln n)$ steps. The advantage of this approach is that we may have a better upper bound if the procedure stops earlier. For example, if G is a complete graph, then the construction stops after 1 step. More generally, for each $x \in S_i \setminus S_{i+1}$, take

a vertex $y \in S_{i+1}$ with $\kappa(x, y) < \kappa_{i+1}$. This is possible since S_{i+1} is a maximal subset. Regarding the pair xy as an edge, this gives a tree with depth at most $O(\ln \ln n)$. Let l be the minimum possible depth. Then the same proof would yield $B = O(l^2 M)$.

6 The sharpness of Main Theorem

In this section we show that both the lower bound and upper bound in the Main Theorem 1.3 are sharp, up to a constant factor. More exactly, we give an example where B and C are of order $\Theta(M)$ and also one where B and C are of order $\Theta(M(\ln \ln n)^2)$.

The proof for the lower bound is easy: for the complete graph on n vertices, all three parameters B, C and M are $\Theta(n \ln n)$.

The construction to match to upper bound is more complicated. It is a tree of depth d defined as follows. The root is at level 1. Each vertex at the i^{th} level has 2^{2^i} children, and the edge between the mother and a child has multiplicity 2^i .

The number of vertices in the i^{th} level is

$$N_i = \prod_{j=1}^{i-1} 2^{2^j} = 2^{2^i - 2}.$$

The number of the vertices in the whole tree is

$$n = \sum_{i=1}^d N_i = \sum_{i=1}^d 2^{2^i - 2} = N_d(1 + o(1)).$$

The number of edges between the i^{th} and $(i+1)^{\text{th}}$ level is $E_i = 2^i N_{i+1} = 2^{2^i + i - 2}$. The total number of edges is

$$E = \sum_{i=1}^{d-1} E_i = E_{d-1}(1 + o(1)).$$

Notice that $d = \Omega(\ln \ln n)$. We first show

$$M = \Theta(E).$$

It is well-known that the commute time between two vertices x, y in a tree (possibly with multiple edge) is

$$2E \cdot \sum_{j=0}^{l-1} \frac{1}{m(x_j x_{j+1})}, \quad (9)$$

where $x = x_0, \dots, x_l = y$ is the path connecting x and y , and $m(vw)$ is the multiplicity of the edge vw . For a lower bound, consider the set S_2 of (four) vertices in level 2. The commute time is $2E$ for any pair (by (9)), which gives $M \geq 2E \ln 4$. For an upper bound, let S be a set of size at least 2. Then take the maximum level i_0 such that there is a vertex in level i_0 having at least two descendants (including itself) in S . Since the multiplicity of an edge geometrically increases as the level increases, (9) implies that the commute time of a pair who has a common ancestor in level i_0 is at most $O(E/2^{i_0})$, especially $\kappa_S = O(E/2^{i_0})$. Moreover, since no pair has a common ancestor in level $i_0 + 1$, the number of vertices in S below level i_0 is at most N_{i_0+1} . Trivially, the number vertices of S above or in level i_0 is at most $\sum_{i=1}^{i_0} N_i = o(N_{i_0+1})$. Thus $|S| = (1 + o(1))N_{i_0+1}$ and

$$\kappa_S \ln |S| = O(E).$$

In the rest of this section, we shall omit floors and ceilings, for the sake of a clearer presentation.

Claim 6.1 *The cover time of this tree satisfies $C = \Omega(Md^2)$.*

Proof. It suffices to show that for a sufficiently large constant K , a walk of length $T = d^2 E/K$, starting from a stationary point, covers the tree with probability at most $1/2$.

To start, set $k = 10 \ln \ln d$ and define a sequence b_i as follows

$$b_k = d^2/\sqrt{K}, b_i = b_{i-1}(1 - \sqrt{\frac{1}{2}/b_{i-1}}),$$

for all $i > k$. Let l be the first index such that $b_l \leq 1/2$. Arithmetic shows that if K is sufficiently large then $l < d - 1$. Set $a_i = 2^i b_i$ and $m_i = 2^{2^{i+1}}$, a simple calculation shows

$$a_i = 2a_{i-1} - \sqrt{\frac{1}{4}a_{i-1} \ln m_{i-1}}.$$

Let X_i denote the minimum number of times a multi-edge from level i to level $i+1$ is crossed in a finite walk. We say that a walk is a T_i -walk if it stops when $X_i = a_i$ and denote by A_i the event that a T_i walk covers the tree. Furthermore, let

B be the event that a walk of length T satisfies $X_k \geq a_k$. Notice that

$$P(\text{A walk of length } T \text{ covers the tree})$$

$$\leq P(B) + P(A_k)$$

The expectation of the number of crosses of any multi-edge between the k^{th} level and the $(k+1)^{th}$ level is $2^k T/E = 2^k d^2/K$, where 2^k is the multiplicity of the edge. On the other hand, $a_k = 2^k d^2/\sqrt{K}$ by definition. Therefore, by Markov's inequality $P(B)$ is at most $1/\sqrt{K} < 1/3$. To finish the proof, we show that $P(A_k) = o(1)$. Observe that for any $i \geq k$, $P(A_i)$ is upper bounded by

$$P(\text{a } T_i\text{-walk satisfies } X_{i+1} \geq a_{i+1}) + P(A_{i+1}).$$

It follows that

$$P(A_k) \leq \sum_{i=k}^{l-1} P(\text{a } T_i\text{-walk satisfies } X_{i+1} \geq a_{i+1}) + P(\text{a } T_l\text{-walk covers the tree}).$$

To show that $P(A_k) = o(1)$, it now suffices to prove that

$$\sum_{i=k}^{l-1} P(\text{a } T_i\text{-walk satisfies } X_{i+1} \geq a_{i+1}) = o(1) \quad (10)$$

and

$$P(\text{a } T_l\text{-walk covers the tree}) = o(1). \quad (11)$$

It will be useful to think about the walk using a “balls and urns” model. Consider a vertex u on level i . Attach to each neighbor of u . Any time we exit node u , drop a ball into the corresponding urn. Then balls will be dropped into the urns independently, so that the urns corresponding to the children of u have the same probability, and the urn corresponding the parent of u has half this probability. Conversely, if for each node, we decide how to drop balls into the urns, then we determine a unique walk. It is important to notice that the number of times an edge is crossed depends only

on the ball distributions corresponding to nodes above the edge.

Assume that the multi-edge between u and its parent v is crossed x times; then the numbers of crossing of the multi-edges going down from u is the same as the number of balls in the big urns at the moment the small urn has x balls.

Using the balls and urns terminology, (10) follows from the following lemma.

Lemma 6.2 *Assume that a and m are large numbers, and $a' = 2a - \sqrt{\frac{1}{4}a \ln m} \geq 0$. Drop balls into one small urn and m big urns until the small urn has a balls, then with probability at least $1 - \left(\exp(-\ln^{2/3} m) + \exp(-m^{1/2})\right)$, one of the big urns has at most a' balls.*

Proof. We use the following fact which is easy to prove. If X is sum of i.i.d. binary random variables and X has large expectation μ , then for any $\sqrt{\mu} \leq L \leq \mu$

$$\exp(-2L^2/\mu) \leq P(X \leq \mu - L) \leq 2\exp(-L^2/2\mu). \quad (12)$$

To prove the lemma, we first show that with probability at least $1 - \exp(-\ln^{2/3} m)$, at the first moment when the small urn has a balls, the number of balls dropped is at most $A = a(2m+1) + 4m\sqrt{a \ln^{2/3} m}$. To show this, it is enough to prove that if one drops A balls randomly into one small urn and m big urns, then with probability at least $1 - \exp(-\ln^{2/3} m)$ the small urn has at least a balls. The number of balls in the small urn can be expressed as a sum of A i.i.d. binary random variables and has expectation $\mu = A/(2m+1) = a + L$, where $L = 2\sqrt{a \ln^{2/3} m} + o(1)$. The claim follows directly from the upper bound in (12), with room to spare.

To finish the proof of the lemma, we show that if we drop A balls into one small urn and m big urns, then there is a big urn with at most a' balls with probability at least $1 - \exp(-m^{1/2+o(1)})$. The number of balls in a fixed big urn is a sum of A i.i.d. binary random variables and has expectation $A/(m+1/2)$. Set $L' = A/(m+1/2) - a'$; it is clear that $L' = (1+o(1))\sqrt{\frac{1}{4}a \ln m}$. We say an urn is

“good” if it has at most a' balls and “bad” otherwise. By the lower bound in (12), the probability that a fixed urn is “good” is at least

$$p = \exp\left(-2L'^2/(A/(m+1/2))\right) \geq m^{-1/2}.$$

So the probability that an urn is “bad” is at most $1-p$. Observe that the events “urn U_1 is bad” and “urn U_2 is bad” are negatively correlated, for any two fixed urns U_1 and U_2 . Using FKG inequality and induction, we can show

$$\begin{aligned} P(\text{all } m \text{ urns are “bad”}) &\leq (1-p)^m \\ &\leq (1-m^{-1/2})^m \leq \exp(-m^{1/2}), \end{aligned}$$

concluding the proof. \square

Now (10) follows from the previous lemma and the fact that

$$\sum_{i=k}^{l-1} \exp(-\ln^{2/3} m_i) + \exp(-m_i^{1/2+o(1)}) = o(1).$$

Here we need to use the condition $k = 10 \ln \ln d$.

To prove (11), it suffices to prove show that if one drops balls into one small urn and $m_l = 2^{l+1}$ big urns until the small urn has $a_l \leq 2^l/2$ balls, then with probability at least $1 - o(1)$, there is an empty big urn. Similar to the proof of Lemma 6.2, one can show that at the time when the small urn has a_l balls, with probability $1 - o(1)$, at most $3a_l m_l$ balls have been dropped (the constant 3 is generous). To conclude, we show that if we drop $A_l = 3a_l m_l$ balls into m_l identical urns, then with probability $1 - o(1)$, there is an empty urn. Since $a_l \leq 2^l/2$ and $m_l = 2^{l+1}$, $A_l \leq \frac{2}{3} m_l \ln m_l$, the claim follows by a standard coupon collector argument. \square

References

- [1] D. J. Aldous and J. Fill, *Time-reversible Markov chains and random walks on graphs* (book in preparation).
- [2] T. Gallai, On directed paths and circuits, in: *Theory of Graphs* (Proc. Colloq., Tihany, 1966) Academic Press, New York (1968), 115–118.

- [3] P. Matthews, Covering problems for Brownian motion on spheres, *Ann. Prob.* **16** (1988), 189–199.
- [4] P. Tetali and P. Winkler, Simultaneous reversible Markov chains, in: *Combinatorics, Paul Erdős is Eighty*, Vol. 1 (ed. D. Miklós, V. T. Sós, T. Szőnyi), Bolyai Society, Budapest, 1993, 422–452.
- [5] P. Winkler and D. Zuckerman, Multiple cover time, *Random Structures Algor.* **9** (1996), 403–411.